arXiv:1601.02974v2 [q-bio.NC] 14 Mar 2016

Memory Recall and Spike Frequency Adaptation

James P. Roach*
Neuroscience Graduate Program, University of Michigan

Leonard M. Sander
Department of Physics, University of Michigan and
Center for the Study of Complex Systems, University of Michigan

Michal R. Zochowski
Department of Physics, University of Michigan
Center for the Study of Complex Systems, University of Michigan and
Biophysics Program, University of Michigan

(Dated: March 15, 2016)

The brain can reproduce memories from partial data; this ability is critical for memory recall. The process of memory recall has been studied using auto-associative networks such as the Hopfield model. This kind of model reliably converges to stored patterns which contain the memory. However, it is unclear how the behavior is controlled by the brain so that after convergence to one configuration, it can proceed with recognition of another one. In the Hopfield model this happens only through unrealistic changes of an effective global temperature that destabilizes all stored configurations. Here we show that spike frequency adaptation (SFA), a common mechanism affecting neuron activation in the brain, can provide state dependent control of pattern retrieval. We demonstrate this in a Hopfield network modified to include SFA, and also in a model network of biophysical neurons. In both cases SFA allows for selective stabilization of attractors with different basins of attraction, and also for temporal dynamics of attractor switching that is not possible in standard auto-associative schemes. The dynamics of our models give a plausible account of different sorts of memory retrieval.

The brain stores memories as patterns of synaptic strengths in the network of neurons. It can store multiple memories and retrieve them in a reliable way, and can change from one to another as attention wanders. However, there is no agreement in the neuroscience community of how this occurs. This paper offers a partial solution to understanding the mechanism for retrieval and switching based on a known physiological effect, spike frequency adapation (SFA).

Decades of work on understanding storage and retrieval have focussed on versions of the Hopfield model (a special form of the Ising model) [1–3]. Hopfield networks have many attractive features: they are auto-associative: that is, memories are recalled from a fragment of their data because the memories are stored in *attractors*, i.e. metastable states.

However, as in any statistical model at zero temperature there is no mechanism for escaping an attractor: a single memory pattern would exist for all time. To overcome this problem an artificial 'temperature' is introduced in Hopfield models to allow switching. This 'temperature' (i.e fast random noise in synaptic current) has no obvious biological origin. Thus, despite the elegance of the model, and its utility in computer science, its application to the brain is problematic. Previous efforts to over come this limitation have used feedback input [4], synaptic depression [5], and adaptive mechanisms [6]. As we will see, SFA allows escape from attractors, and, in some cases acts in the same way as the Hopfield temperature. In addition, it turns off particular memories rather than globally smearing out all of them, as temperature does.

SFA is an activity induced reduction in neural firing rate induced by a hyperpolarizing current that grows as a neuron fires – neurons that fire a good deal tend to stop firing after a delay. Thus SFA is a natural mechanism to turn off activity. Further, SFA can be controlled by neuromodulators such as Acetylcholine (ACh), an important regulator of neural excitability. ACh causes a reduction in SFA and provides for its dynamic regulation [7, 8]. In previous work [9] we have presented a network model of Hodgkin-Huxley (HH) neurons with SFA and 'Mexican Hat' coupling which reproduces many features of cortical activity as ACh levels change between sleep and waking states. In the present work (below) we use this model to concentrate on memory retrieval. However, we claim that the essentials of our results are quite robust and independent of the details of the neuron model. To show this, we first consider a version of the Hopfield model [1–3] which has SFA.

In our model we consider networks composed of N=1000 spins, $S = \{s_i\}$ where $s_i = \pm 1$. The network is fully connected with weights $\sigma_{i,j}$. As usual, spin up corresponds to a neuron that fires, and spin down to a silent one. Each spin gets an input:

$$h_i(t) = \sum_{j=1}^{N} \sigma_{i,j} s_j - \theta_i(t), \tag{1}$$

where $\theta_i(t)$ is a local offset field at site i which changes

slowly in time. The first term is the usual Hopfield-Ising term and the second represents SFA.

The dynamics of the spins are as follows: at each time step a random spin is flipped with probability:

$$P_h(s_i) = \frac{1}{1 + e^{-2s_i h_i/T}}, \tag{2}$$

where T is the noise. In much of what follows we take T to be very small so that P_h is essentially a step function.

The dynamics of θ_i is:

$$\theta_i(s_i) = \frac{A}{1 + e^{-s_i(\hat{t}-\tau_1)/\tau_2}}. \tag{3}$$

Here, $\hat{t}$ is the time since the last state change of the spin, and $\tau_{1,2}$ are time constants which govern the dynamics of attractors. The field θ increases to A for up spins and decreases to zero for down spins. The time constant τ_1 is the time to the half-maximum value of θ_i. We take $\tau_1 = 5$ (timesteps/N), except for the data in Figure 1 where $\tau_1 = 1.5$ (timesteps/N). The rate at which SFA activates/ deactivates is controlled by τ_2 which is set to 0.2, except for the data in Figure 4 where $\tau_2 = 0.6$ (timesteps/N). This implementations of adaptation is different than others in the hopfield model [6]. because it integrates over a longer time (i.e. considers more than the activity at the previous time step). This more closely resembles adaptation in biophysical models.

In the Hopfield scheme memories are stored as attractors, i.e. metastable configurations, $\Xi^\mu = \{\xi_i^\mu\}$. We encode attractors using a modified Hebbs rule [1]:

$$\sigma_{i,j} = \frac{1}{NW} \sum_\mu^p w_\mu \xi_i^\mu \xi_j^\mu. \tag{4}$$

Each attractor is given a weight, w_μ and $W = \sum_\mu^p w_\mu$. Thus $\sigma_{i,j} = 1$ for two spins with correlated activity across all attractors, Ξ^μ, and -1 for spins with anti-correlated activity. We set $w_1/w_{p-1} = 0.5$ and $w_p = 1$, except for the data reported in Figure 4D where all the weights $w_i = 1.0$. The saturation is defined as $\alpha = p/N$. Attractors encoded with lower w_μ are weaker attractors.

In order to determine if the dynamics has settled into the various Ξ^μ we measure the overlap between the stored memory and the current state:

$$m_\mu = \frac{1}{N} \sum_i^N s_i \xi_i^\mu, \tag{5}$$

which is ± 1 when $S = \pm\Xi$ and 0 when $S \perp \Xi$. In each simulation S was always initialized to a random weak attractor, Ξ_{weak}.

In the usual Hopfield model the preference for local, global, or no attractors changes as the noise, T, increases [1, 10]. This transition depends on the storage capacity α [1]. We find very similar behavior as we increase the magnitude of SFA, i.e. A. To show this we compare the standard T versus α plot with a plot of A versus α in Figure 1. The right panel shows how T and α interact to affect the stability of the strong and weak attractors and, eventually, to destabilize all attractors. For small T the dynamics keeps the system in a weak attractor (black on the colormap); for larger T the system enters a regime of stability of stronger attractors (white). For large T no attractors are stable (gray).

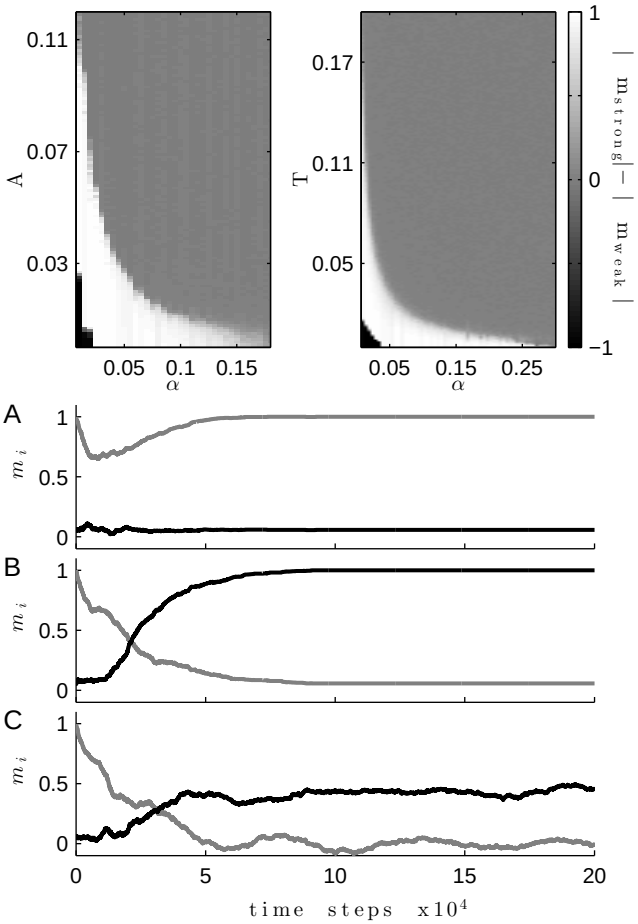

FIG. 1. SFA and T control attractor stability in Hopfield networks. With noise Hopfield networks have three functional states: stability of local attractors, stability of global attractors, and stability of no attractor. The stability of local versus global versus no attractors is shown by the ability of the attractor to move from a weak attractor to a strong one, which is quantified by the difference of the m_μ of the strong attractor and the weak attractor in which the system was initialized. (*Top*)The saturation of memories (α), noise (T), and adaptation (A) affect stability in a similar manner. For low levels of noise local attractors are stable (black). For a given α either increasing A (left) or T (right) leads to a strong attractors being stable (white). Further increase destabilizes all attractors (gray). (*Bottom*) Example dynamics of memory overlap, m_i, for strongly (black) and weakly (gray) weighted memories. In each case $\alpha = 0.01$; adaptation levels are (A) A = 0.01 (B) A = 0.05 (C) A = 0.3.

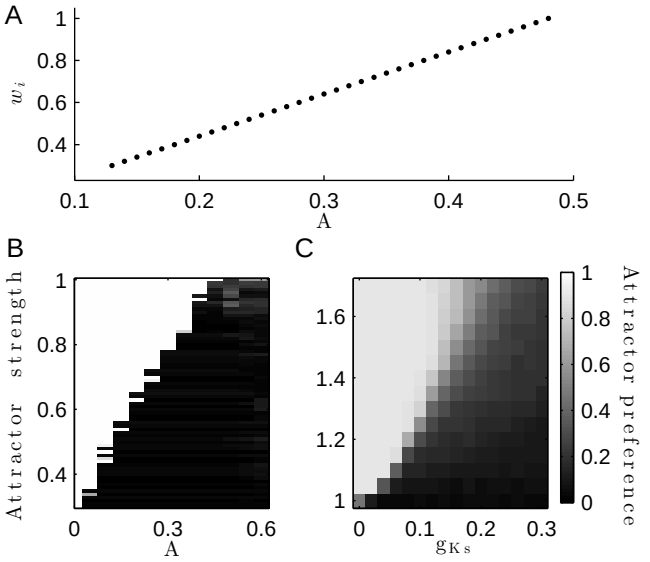

FIG. 2. Attractor stability varies as a function of A. For SFA to induce the network to leave a attractor it must be large enough to overcome the energy barrier of the attractor. Mean field calculations predict a linear relationship between the strength of an attractor and the amount of adaptation, A, to destabilize it (A).This is best seen in the Hopfield model (B). The threshold value of A increases linearly as the attractor strength, w_μ increases. A similar effect is seen in the spiking network model (C).

SFA has a very similar effect; see Figure 1, left. To see how this comes about, consider the Hamiltonian:

$$E = -\frac{1}{2}\sum_{i,j}^{N}\sigma_{i,j}s_i s_j + \sum_{i}^{N}\theta_i(t)s_i, \tag{6}$$

which defines a slowly varying energy landscape. The first term is the ordinary Ising energy, and the second can be thought of as a time-dependent magnetic field that increases (for up spins) to A when $t \gg \tau_2$. Increasing A destabilizes minima in E; an attractor becomes unstable for $t \to \infty$ if A is large enough. In Figure 2 we show the magnitude of SFA required to cause the system to leave an attractor of a given strength. It increases linearly as attractor strength increases. A similar effect occurs in the HH model, below.

The stability of an attractor of a given weight can also be investigated by mean field theory [10, 11]. The mean field equations for the system are:

$$\langle s_i\rangle = \tanh(\frac{\beta}{NW}\sum_{j,\mu} w_\mu \xi_i^\mu \xi_j^\mu \langle s_j\rangle - 2\theta_i), \tag{7}$$

where $\beta = 1/T$. By exploiting the fact that $\langle s_i\rangle = m\xi_i^\upsilon$ the mean field equations can be rewritten as:

$$m\xi_i^\upsilon = \tanh(\frac{\beta}{NW}\sum_{j,\mu} w_\mu \xi_i^\mu \xi_j^\mu m\xi_i^\upsilon - 2\theta_i). \tag{8}$$

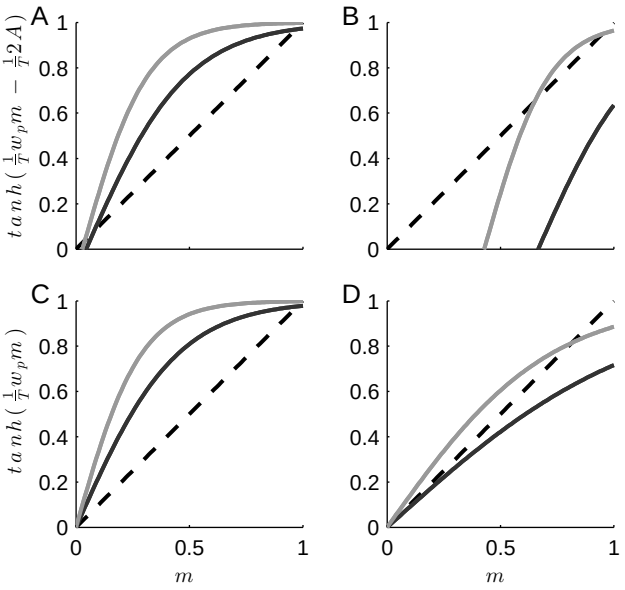

FIG. 3. SFA and T destabilize attractors of different strengths in Hopfield networks. Solutions to mean field equations illustrate how adaptation and temperature destabilize weak attractors. When A is low many memories are stabile (panel A; A=0.1). Increasing A destabilizes weak memories while preserving strong memories (panel B; A=0.25). When adaptation is absent temperature has a similar effect where all memories are stable for low T (panel C; $T = 0.2$), while only strong memories are stable for high T (panel D; $T = 0.5$). The dashed line shows $m = m$; the solid black line shows the mean field equation for $w_\upsilon = 0.45$; the solid gray line shows the mean field equation for $w_\upsilon = 0.75$.

If $p \ll N$ any overlap between memories is negligible so the mean field equation and the system is in memory υ for a time $\gg \tau_1$ becomes:

$$m\xi_i^\upsilon = \tanh(\beta w_\upsilon m\xi_i^\upsilon - \frac{1}{T}2A), \tag{9}$$

which be simplified to $m = \tanh(\beta w_\upsilon m - \beta 2A)$. When this equation has solutions beyond $m = 0$ a memory with strength w_υ is stable for a given T or A. Figure 3 shows mean field solutions for memories with strengths $w_\upsilon = 0.45$ (solid black line) and $w_\upsilon = 0.75$ (solid gray line). As in the numerical results adaptation (3 top panels) and temperature (bottom panels) have similar effects on the stability of memories. For low levels ($A = 0.1$, $T = 0.2$; left panels) both strong and weak memories are stable (i.e. both have solutions beyond $m = 0$), but moderate increases in A or T destabilize weaker memories ($A = .25$, $T = 0.5$; right panels)

Thus, changes in the strength of SFA can play the same role as changes in T by destabilizing attractors of varying strength as A increases. Interesting time-dependent effects occur for intermediate values of A when the τ's are not too large. Because θ_i is a function of t we can

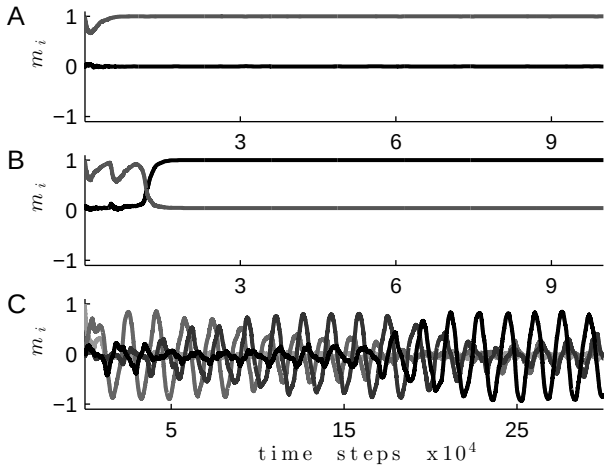

FIG. 4. Examples of network dynamics in the modified Hopfield model. For small A even a weak attractor is stable (panel A; A = 0.01). A moderate increase leads to the strong attractor becoming stable (B, A = 0.1). For larger A damped oscillations with period $\sim 4\tau_1$ emerge: C, A = 0.4. In panels A and B the gray line corresponds to a weak attractor and the black corresponds to a strong one. In panel C, the lightest gray is the weakest attractor. All other lines are strong attractors of equal weight.

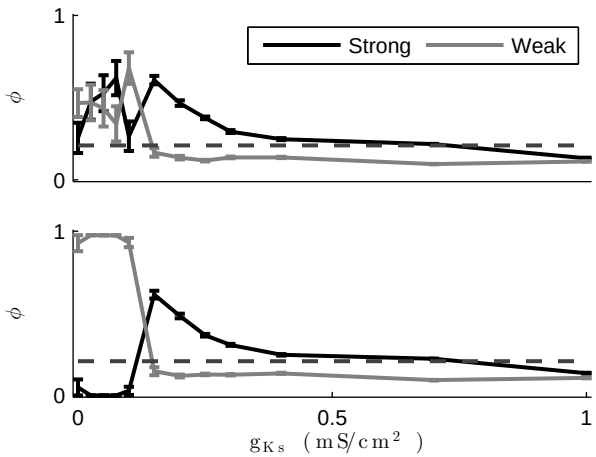

FIG. 5. Attractor preference and g_{Ks}. The quantity ϕ is the fraction of time that activity is located within an attractor. Black dashed line, control value. (Top) For initial locations outside any attractor no clear preference emerges for small g_{Ks}. For moderate g_{Ks} there is clear preference for the strong attractor. (Bottom) For initial conditions within the weak attractor activity never leaves for small g_{Ks}. There is significant preference for the strong attractor for moderate g_{Ks}.

generate chains of attractor preferences, as opposed to stability in a deep attractor or a random walk (as in the standard Hopfield model for large T). These results are shown in Figure 4. For small A local, weak attractors are stable. A moderate increase leads to strong attractors being stable (Figure 4 A, B). Further increase of A leads to oscillations of period $\sim 4\tau_1$; Figure 4 C. This is similar to the latching dynamics found in [6].

We next turn to a more realistic neuron model to compare the effects of SFA for the two cases. The spiking network model introduced previously [9] considers $N_E = 1225$ excitatory and $N_I = 324$ inhibitory HH neurons arrayed on two square lattices of size $L_{E/I}$. The coupling was of lateral inhibition (Mexican Hat) type where short range excitation is balanced with global inhibition. All excitatory neurons were connected to neighbors within radius $R_{xx} = \sqrt{L^2_{E/I} k_{xx}/\pi N_{E/I}}$ where $k_{ei} = 16$ is the degree of excitatory to excitatory connections, $k_{ei} = 4$ is the degree of excitatory to inhibitory connections. Neural dynamics were modeled by the current balance equation [12]:

$$c_m \frac{dV_i}{dt} = -g_{Na} m_\infty^3 h(V_i - E_{Na}) - g_{Kdir} n^4 (V_i - E_K) - g_{Ks} s (V_i - E_K) - g_L (V_i - E_L) - I_{syn,i} + I_{ext} \quad (10)$$

In this equation, as we will see, g_{Ks} sets the magnitude of the SFA; it corresponds to A in the model above.

The dynamics of the gating variables h, n and s is of the form $dx/dt = (x_\infty(V) - x)/\tau_x(V)$ with additional specific evolution of the two voltage dependent parameters x_∞ and τ_x. The slow potassium current conductance, g_{Ks} controls the level of SFA (i.e. lower values of g_{Ks} correspond to low SFA). The level of ACh modulates g_{Ks}: the maximum ($g_{Ks} = 1.5$ mS/cm^2) and minimum ($g_{Ks} = 0$) correspond to the absence or maximum of ACh, respectively. For more details see [9].

The synaptic current to neuron i is $I_{syn,i} = g_E(t)(V_i - E_E) + g_I(t)(V_i - E_I)$ and the dynamics of $g_{E/I}(t)$ is:

$$g_{E/I}(t) = K \sum_{j}^{\in E/I} \sigma_{i,j} (e^{\frac{-(\tilde{t}_j - \tau_D)}{\tau_S}} - e^{\frac{-(\tilde{t}_j - \tau_D)}{\tau_F}}) \quad (11)$$

where $\sigma_{i,j}$ is the synaptic weight between neuron i and neuron j, and $\tau_{S,F}$ are time constants equal to 3.0 and 0.3 ms respectively. $\sigma_{i,j}$ is set to 0.02 mS/cm^2 unless otherwise stated, $\tilde{t}_j$ is the time since the last spike of neuron j, and K is a normalization constant. I_{ext} set so that all neurons fire at 10 Hz in the absence of synaptic input and g_{Ks}. The equations were integrated using the 4th order Runge-Kutta method at a 0.05 ms time step to 20 s. Data points are averages of 20 sets of initial conditions.

We have shown [9] that the nature of the dynamics is that for small g_{Ks} (large concentrations of ACh) there is a stationary, localized region (a 'bump') of spiking activity. For larger g_{Ks} the bump travels through the lattice.

To consider memory we introduce spatial attractors by increasing synaptic strength in certain locations [9, 13]. These attractors fix the location of the bump when the

dynamics is in the stationary regime. For a single attractor, preference for the attractor falls as g_{Ks} increases [9]. This model is quite different from the one discussed above: the excitatory coupling is short-ranged in contrast to the Hopfield σ's which are long-ranged. The attractors here are defined by local geometry. Nevertheless, SFA gives common results for the two cases.

To consider multiple attractors of variable strength we strengthened synaptic strength in two network regions at opposite ends of the lattice. The strong attractor had 100% stronger excitatory connections and the weak attractor had a 50% increase. To examine how network preference changed as a function of g_{Ks} in multi-attractor networks we did simulations where activity was initialized by injecting a 0.25 μA/cm^2 current to a region outside either of the attractors for the first 0.5 s of the simulation. Preference for a given attractor was quantified by the measure ϕ which is the fraction of time that the center of the bump, calculated according to [14], is located within the attractor. For low levels of SFA there is no clear preference indicating that activity localizes to attractors randomly; Figure 5, top. Increasing g_{Ks} leads to a clear preference for the strong attractor. The preference for any attractor disappears for large g_{Ks} as the network enters the regime of traveling bump dynamics.

To further test the stability of the two attractors we put activity initially on the weak attractor. For small g_{Ks} activity remained localized there: Figure 5, bottom. For larger g_{Ks} activity moved to the stronger attractor. This confirms that the strength of SFA can control the stability of attractors having different depths.

We also considered how the relative stability of the attractors depends on the ratio of inhibitory to excitatory coupling, $w_{e/i}$). Figure 6 is a phase plot of final attractor preference for different $w_{e/i}$) and g_{Ks}. The relative preference for the two attractors was measured by $\phi_{strong} - \phi_{weak}$, which ranges between 1 (preference for the strong attractor) and -1 (preference for the weak). Interestingly, weakening inhibition abolishes any preference for the strong attractor at intermediate levels of g_{Ks}.

Thus, for both models SFA can selectively destabilize attractors effectively controlling attractor preference. With SFA we can have long-term preferential activation of attractors of different strengths and also non-trivial time dependence of attractor sequences. This provides a biologically plausible mechanism for switching between encoded patterns [10, 15].

In the biophysical model SFA depends on ACh [8, 12, 16]. Our results imply that ACh controls memory retrieval dynamics. Note the relevance of our results to context dependent release of ACh and its role in attention [17]. During tasks requiring a high degree of focus, low SFA allows the brain to fix on the memory that closely fits the current sensory input. On the other hand, with high ACh the attractor can be reinforced by synaptic plasticity. As attention requirements are relaxed, and

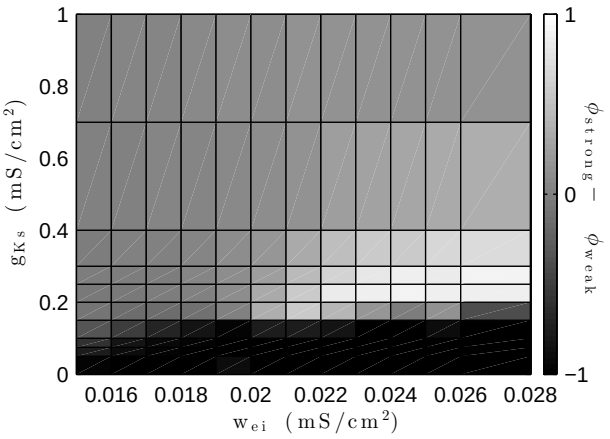

FIG. 6. Strong attractor preference is controlled by inhibition strength in spiking networks. We measure the differential attractor preference for the two attractors by ($\phi_{strong} - \phi_{weak}$), which is 1 when all activity is located within the strong attractor and -1 when all activity is within the weak. The preference for the strong attractor at moderate SFA disappears when inhibition is decreased.

ACh levels fall, moderate levels of SFA allow for sampling of the memory space, see [18].

The largest variation in cortical ACh levels occurs between sleep/ wake states. In this case the highest levels occur during rapid eye movement (REM) sleep and the lowest during slow wave sleep [19]. We argue that intermediate levels of ACh during wake states allow for memory recall when externally driven network states are allowed to wander to find the optimal state. REM sleep is thought to be important for memory consolidation, where retrieval of weakly stored attractors of previous experience is essential to their consolidation. NREM sleep associated with low ACh levels is characterized by slow waves and may play a role in synaptic rescaling [20].

JPR was supported by an NSF Graduate Research Fellowship Program under Grant No. DGE 1256260 and a UM Rackham Merit Fellowship. MRZ and LMS were supported by NSF PoLS 1058034.

* roachjp@umich.edu

[1] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Annals of Physics **173**, 30 (1987).
[2] M. E. Hasselmo, B. P. Anderson, and J. M. Bower, Journal of Neurophysiology **67**, 1230 (1992).
[3] J. J. Hopfield, Proceedings of the National Academy of Sciences **79**, 2554 (1982).
[4] S. Recanatesi, M. Katkov, S. Romani, and M. Tsodyks, Frontiers in computational neuroscience **9**, 275 (2015).
[5] I. Lerner and O. Shriki, Frontiers in psychology **5**, 314 (2014).
[6] A. Akrami, E. Russo, and A. Treves, Brain research

1434, 4 (2012).
[7] S. P. Aiken, B. J. Lampe, P. A. Murphy, and B. S. Brown, British journal of pharmacology **115**, 1163 (1995).
[8] A. C. Tang, A. M. Bartels, and T. J. Sejnowski, Cerebral Cortex **7**, 502 (1997).
[9] J. P. Roach, E. Ben-Jacob, L. M. Sander, and M. Zochowski, PLoS Computational Biology **11**, e1004449 (2015).
[10] M. Lewenstein and A. Nowak, Physical Review Letters **62**, 225 (1989).
[11] D.J. Amit, H. Gutfreund, and H. Sompolinsky, Physical review. A **32**, 1007 (1985).
[12] K. M. Stiefel, B. S. Gutkin, and T. J. Sejnowski, Journal of Computational Neuroscience **26**, 289 (2008).
[13] A. Renart, P. Song, and X.-J. Wang, Neuron **38**, 473 (2003).
[14] L. Bai and D. Breen, Journal of Graphics, GPU, and Game Tools **13**, 53 (2008).
[15] R. Monasson and S. Rosay, Physical Review Letters **115**, 098101 (2015).
[16] Y. Tsuno, N. W. Schultheiss, and M. E. Hasselmo, The Journal of Physiology **591**, 2611 (2013).
[17] M. E. Hasselmo and M. Sarter, Neuropsychopharmacology **36**, 52 (2011).
[18] J. J. Hopfield, Proceedings of the National Academy of Sciences **107**, 1648 (2010).
[19] J. Vazquez and H. A. Baghdoyan, American journal of physiology. Regulatory, integrative and comparative physiology **280**, R598 (2001).
[20] G. Tononi and C. Cirelli, Brain Research Bulletin **62**, 143 (2003).